

IN THE CLAIMS:

- 1 1. (Currently Amended) A method for uniformly distributing data transmitted by a server
2 over a plurality of underlying links of an aggregate within a computer network, the
3 method comprising the steps of:
 - 4 defining a unit of data as a datagram;
 - 5 apportioning each datagram into at least one fragment at the server;
 - 6 associating each fragment to an underlying link of the aggregate on the basis of an
7 Internet protocol (IP) identifier (ID) of each datagram and a number of active links of the
8 aggregate; and
 - 9 transmitting the fragment over its associated underlying link from the server to the
10 computer network to transmit fragments of the datagram over the same associated under-
11 lying link.
- 1 2. (Original) The method of Claim 1 wherein the step of associating comprises the step of
2 producing a result representing a remainder upon dividing the IP ID by the number of
3 active links.

- 1 3. (Original) The method of Claim 2 wherein the step of associating further comprises the
2 steps of:
 - 3 calculating the IP ID of each datagram in a sequential manner; and
 - 4 rotating the fragments of each datagram among all the underlying links to thereby
5 ensure that all fragments having the same IP ID are provided to the same physical link of
6 the aggregate.

1 4. (Original) The method of Claim 1 wherein the step of associating comprises the steps
2 of:

3 logically combining the IP ID with a predetermined mask to produce a quantity;
4 right shifting the quantity a predetermined number of places; and
5 establishing a threshold at which a group of data is forwarded to each underlying
6 link of the aggregate.

1 5. (Original) The method of Claim 4 wherein the step of associating further comprises the
2 step of producing a result representing a remainder upon dividing the right shifted logi-
3 cally combined quantity IP ID and predetermined mask by the number of active links.

1 6. (Currently Amended) The method of Claim 5 A method for uniformly distributing data
2 transmitted by a server over a plurality of underlying links of an aggregate within a com-
3 puter network, the method comprising the steps of:

4 defining a unit of data as a datagram;
5 apportioning each datagram into at least one fragment at the server;
6 associating each fragment to an underlying link of the aggregate on the basis of an
7 Internet protocol (IP) identifier (ID) of each datagram and a number of active links of the
8 aggregate, wherein the step of associating includes
9 logically combining the IP ID with a predetermined mask to produce a
10 quantity,
11 right shifting the quantity a predetermined number of places,
12 establishing a threshold at which a group of data is forwarded to each un-
13 derlying link of the aggregate, and

14 producing a result representing a remainder upon dividing the right shifted
15 logically combined quantity IP ID and predetermined mask by the number of ac-
16 tive links, wherein the IP ID is a 16-bit value, the predetermined mask is 0xFF80
17 and predetermined number of right shifted places is 7, and wherein the group of
18 data comprises 128 IP IDs;
19 transmitting the fragment over its associated underlying link from the server to the
20 computer network.

1 7. (Original) The method of Claim 6 wherein the group of data comprises one of 128 dif-
2 ferent transport control protocol (TCP) fragments and 128 different user datagram proto-
3 col (UDP) datagrams.

1 8. (Original) The method of Claim 7 wherein each UDP datagram comprises up to 23
2 fragments.

1 9. (Original) The method of Claim 1 further comprising the steps of:
2 loading at least one data buffer of the server with the at least one fragment;
3 fetching the fragment from the data buffer; and
4 loading at least one queue of the server with the fragment, the queue associated
5 with the underlying link.

1 10. (Original) A system adapted to uniformly distributing data over a plurality of underly-
2 ing links of an aggregate within a computer network, the system comprising:
3 a processor;

4 a memory coupled to the processor and having locations addressable by the proc-
5 essor;

6 an operating system resident in the memory locations and executed by the proces-
7 sor, the operating system configured to implement a modified load balancing technique
8 that defines a unit of data as a datagram, the operating system comprising an Internet Pro-
9 tocol (IP) layer that apportions the datagram into at least one fragment, the operating sys-
10 tem further comprising a virtual interface process that associates the fragment to an un-
11 derlying link of the aggregate on the basis of an IP identifier (ID) of the datagram and a
12 number of active links of the aggregate; and

13 at least one network adapter coupled to the memory and processor that cooperates
14 with a network driver of the operating system to transmit the fragment over the associated
15 underlying link to the computer network.

1 11. (Currently Amended) Apparatus for uniformly distributing data transmitted by a
2 server over a plurality of underlying links of an aggregate within a computer network, the
3 apparatus comprising:

4 means for defining a unit of data as a datagram;
5 means for apportioning each datagram into at least one fragment at the server;
6 means for associating each fragment to an underlying link of the aggregate on the
7 basis of an Internet protocol (IP) identifier (ID) of each datagram and a number of active
8 links of the aggregate; and

9 means for transmitting the fragment over its associated underlying link from the
10 server to the computer network to transmit fragments of the datagram over the same as-
11 sociated underlying link.

1 12. (Original) The apparatus of Claim 11 wherein the means for associating comprises
2 means for producing a result representing a remainder upon dividing the IP ID by the
3 number of active links.

1 13. (Original) The apparatus of Claim 12 wherein the means for associating further com-
2 prises:

3 means for calculating the IP ID of each datagram in a sequential manner; and
4 means for rotating the fragments of each datagram among all the underlying links
5 to thereby ensure that all fragments having the same IP ID are provided to the same
6 physical link of the aggregate.

1 14. (Original) The apparatus of Claim 11 wherein the means for associating comprises:

2 means for logically combining the IP ID with a predetermined mask to produce a
3 quantity;

4 means for right shifting the quantity a predetermined number of places; and
5 means for establishing a threshold at which a group of data is forwarded to each
6 underlying link of the aggregate.

1 15. (Original) The apparatus of Claim 14 wherein the means for associating further com-
2 prises means for producing a result representing a remainder upon dividing the right
3 shifted logically combined quantity IP ID and predetermined mask by the number of ac-
4 tive links.

1 16. (Original) A computer readable medium containing executable program instructions
2 for uniformly distributing data transmitted by a server over a plurality of underlying links

3 of an aggregate within a computer network, the executable program instructions comprising program instructions for:

5 defining a unit of data as a datagram;

6 apportioning each datagram into at least one fragment at the server;

7 associating each fragment to an underlying link of the aggregate on the basis of an
8 Internet protocol (IP) identifier (ID) of each datagram and a number of active links of the
9 aggregate; and

10 transmitting the fragment over its associated underlying link from the server to the
11 computer network.

1 17. (Original) The computer readable medium of Claim 16 wherein the program instruction
2 for associating comprises a program instruction for producing a result representing a
3 remainder upon dividing the IP ID by the number of active links.

1 18. (Original) The computer readable medium of Claim 17 wherein the program instruction
2 for associating further comprises program instructions for:

3 calculating the IP ID of each datagram in a sequential manner; and

4 rotating the fragments of each datagram among all the underlying links to thereby
5 ensure that all fragments having the same IP ID are provided to the same physical link of
6 the aggregate.

1 19. (Original) The computer readable medium of Claim 16 wherein the program instruction
2 for associating comprises program instructions for:

3 logically combining the IP ID with a predetermined mask to produce a quantity;

4 right shifting the quantity a predetermined number of places; and

5 establishing a threshold at which a group of data is forwarded to each underlying link of
6 the aggregate.

1 20. (Original) The computer readable medium of Claim 19 wherein the program instruc-
2 tion for associating further comprises the program instruction for producing a result rep-
3 resenting a remainder upon dividing the right shifted logically combined quantity IP ID
4 and predetermined mask by the number of active links

1 21. (Currently Amended) A method for distributing data over a plurality of network links
2 within a computer network, comprising the steps of:

3 defining a unit of data as a datagram;

4 apportioning each datagram into at least one fragment;

5 | associating each fragment datagram to a network link of the plurality of network
6 links according to a round robin policy based at least in part on an Internet protocol (IP)
7 identifier (ID) of each datagram;

8 | transmitting the fragments of the datagram over the fragment's datagram's associ-
9 ated network link.

1 22. (Previously Presented) The method of claim 21 wherein the step of associating is
2 further based, at least in part, on a number of network links in the plurality of links.

1 23. (Previously Presented) The method of claim 22 wherein the step of associating com-
2 prises the step of:

3 producing a result representing a remainder by dividing the IP ID by the number
4 of network links.

1 24. (Previously Presented) The method of claim 23 wherein the step of associating fur-
2 ther comprises the steps of:

3 calculating the IP ID of each datagram in a sequential manner; and
4 rotating the fragments of each datagram among all the network links of the plural-
5 ity of network links to thereby ensure that all fragments having the same IP ID are pro-
6 vided to the same network link.

1 25. (Previously Presented) The method of claim 21 wherein the step of associating com-
2 prises the steps of:

3 logically combining the IP ID with a predetermined mask to produce a quantity;
4 right shifting the quantity by a predetermined number of places to create a new
5 quantity; and
6 establishing a threshold at which a group of data is forwarded to each network
7 link of the plurality of network links.

1 26. (Previously Presented) The method of claim 25 wherein the step of associating fur-
2 ther comprises the step of producing a result representing a remainder upon dividing the
3 new quantity by a number of network links in the plurality of network links.

1 27. (Currently Amended) A system for distributing data over a plurality of network links
2 within a computer network comprising:

3 a processor;
4 a memory coupled to the processor and having locations addressable by the proc-
5 essor;

6 an operating system resident in the memory locations and executed by the processor,
7 the operating system configured to implement a load balancing technique that defines
8 a unit of data as a datagram, the operating system further configured to include an Internet
9 Protocol (IP) layer that apportions the datagram into at least one fragment, the operating
10 system further configured to include a virtual interface process that associates each

11 fragment datagram to a network link of the plurality of network links according to a
12 round robin policy based at least in part on an Internet protocol (IP) identifier (ID) of
13 each datagram; and

14 at least one network adapter coupled to the memory and processor to transmit the
15 fragments of the datagram over the fragment's datagram's associated network link.

1 28. (Currently Amended) A system for distributing data over a plurality of network links
2 within a computer network, comprising:

3 means for defining a unit of data as a datagram;

4 means for apportioning each datagram into at least one fragment;

5 means for associating each fragment datagram to a network link of the plurality of
6 network links according to a round robin policy based at least in part on an Internet pro-
7 tocol (IP) identifier (ID) of each datagram;

8 means for transmitting the fragments of the datagram over the fragment's data-
9 gram's associated network link.

1 29. (Currently Amended) A computer readable medium containing executable program
2 instructions for execution on a processor, the executable program instructions comprising
3 program instructions for:

4 defining a unit of data as a datagram;

5 apportioning each datagram into at least one fragment;

6 | associating each fragment datagram to a network link of the plurality of network
7 | links according to a round robin policy based at least in part on an Internet protocol (IP)
8 | identifier (ID) of each datagram;
9 | transmitting the fragments of the datagram over the fragment's datagram's associ-
10 | ated network link.

1 | 30. (Previously Presented) A method for distributing data over a plurality of network
2 | links within a computer network, comprising the steps of:

3 | dividing a first datagram, having a first Internet protocol (IP) identifier (ID), into
4 | one or more fragments, each fragment of the first datagram associated with the first
5 | Internet protocol (IP) identifier (ID);

6 | selecting a first network link of the plurality of network links for transmission of
7 | the one or more fragments of the first datagram;

8 | transmitting all of the one or more fragments associated with the first IP ID over
9 | the first network link;

10 | dividing a second datagram, having a second Internet protocol (IP) identifier (ID),
11 | into one or more fragments, each fragment of the second datagram associated with the
12 | second Internet protocol (IP) identifier (ID);

13 | selecting a second network link of the plurality of network links for transmission
14 | of the one or more fragments of the second datagram; and

15 | transmitting all of the one or more fragments associated with the second IP ID
16 | over the second network link.

1 | 31. (Previously Presented) The method of claim 30 wherein the first network link and
2 | the second network link are selected according to a round robin policy based on the IP ID
3 | of each datagram.

1 32. (Previously Presented) A system for distributing data over a plurality of network
2 links within a computer network comprising:
3 a processor;
4 a memory coupled to the processor and having locations accessible by the proces-
5 sor;
6 an operating system resident in the memory and executed by the processor, the
7 operating system configured to implement a load balancing technique that divides a first
8 datagram, having a first Internet protocol (IP) identifier (ID), into one or more fragments,
9 each fragment of the first datagram associated with the first Internet protocol (IP) identi-
10 fier (ID), select a first network link of the plurality of network links for transmission of
11 the one or more fragments of the first datagram, transmit all of the one or more fragments
12 associated with the first IP ID over the first network link, divide a second datagram, hav-
13 ing a second Internet protocol (IP) identifier (ID), into one or more fragments, each frag-
14 ment of the second datagram associated with a second Internet protocol (IP) identifier
15 (ID), select a second network link of the plurality of network links for transmission of the
16 one or more fragments of the second datagram, and transmit all of the one or more frag-
17 ments associated with the second IP ID over the second network link.

1 33. (Previously Presented) The system of claim 32 wherein the first network link and
2 the second network link are selected according to a round robin policy based on the IP ID
3 of each datagram.

1 34. (Currently Amended) The method of claim 1 wherein the step of associating further
2 comprises apportioning data ~~substantially~~ equally over the plurality of underlying links of
3 the aggregate within the computer network.

1 35. (Currently Amended) The system of claim 10 wherein the virtual interface process is
2 | configured to apportion data ~~substantially~~ equally over the plurality of underlying links of
3 | the aggregate within the computer network.

1 36. (Currently Amended) The system of claim 11 wherein the means for associating is
2 | configured to apportion data ~~substantially~~ equally over the plurality of underlying links of
3 | the aggregate within the computer network.

1 37. (Previously Presented) The system of claim 27 wherein the round robin policy is
2 further based, at least in part, on a number of network links in the plurality of links.

1 38. (Previously Presented) The system of claim 27 wherein the virtual interface process
2 is configured to produce a result representing a remainder by dividing the IP ID by the
3 number of network links.

1 39. (Previously Presented) The method of claim 38 wherein the virtual interface process
2 is further configured to calculate the IP ID of each datagram in a sequential manner and
3 to rotate the fragments of each datagram among all the network links of the plurality of
4 network links to thereby ensure that all fragments having the same IP ID are provided to
5 the same network link.

1 40. (Previously Presented) The system of claim 27 wherein the virtual interface process
2 is configured to logically combine the IP ID with a predetermined mask to produce a
3 quantity, right shift the quantity by a predetermined number of places to create a new

- 4 quantity, and establish a threshold at which a group of data is forwarded to each network
- 5 link of the plurality of network links.

- 1 41. (Previously Presented) The system of claim 40 wherein the virtual interface process
- 2 is further configured to produce a result representing a remainder upon dividing the new
- 3 quantity by a number of network links in the plurality of network links.

1 Please add new claims 42 *et al.*

1 42. (New) The method of claim 30, wherein the steps of transmitting are performed in
2 parallel.

1 43. (New) A method for uniformly distributing data transmitted by a server over a plural-
2 ity of underlying links within a computer network, comprising:

3 defining a plurality of datagram with each datagram defined as a unit of data;

4 apportioning each datagram into one or more fragments;

5 using an internet protocol (IP) identifier of each datagram to assign a link of the
6 plurality of links to each datagram;

7 associating each fragment of a datagram with an IP identifier of the datagram;

8 and

9 transmitting the fragments of the datagram through the assigned link with the
10 same IP identifier.